

What is claimed is:

1. A method for cluster management in a network environment, comprising:
 - receiving one or more traffic measurements from one or more nodes associated with one or more clusters in said network environment;
 - determining, in accordance with the received traffic measurements, one or more reclustering operations to be performed in said network environment; and
 - dispatching data to realize said reclustering.
2. The method of claim 1, wherein said network environment is a peer-to-peer environment.
3. The method of claim 1, wherein said reclustering operations comprise creation of a new cluster.
4. The method of claim 1, wherein said reclustering operations comprise elimination of one of said clusters.
5. The method of claim 1, wherein said reclustering operations comprise transfer of one or more of said nodes among between one or more of said clusters.
6. The method of claim 1, wherein said traffic measurements are constantly taken.
7. The method of claim 1, wherein said traffic measurements are taken in response to a request for said measurements.

8. The method of claim 7, wherein said measurements are taken for a specified period of time.
9. The method of claim 1, wherein said traffic measurements comprise measurements corresponding to node index updates.
10. The method of claim 1, wherein said traffic measurements comprise measurements corresponding to entity index updates.
11. The method of claim 1, wherein said traffic measurements comprise measurements corresponding to entity queries.
12. The method of claim 1, wherein determining comprises determination of ideal cluster size.
13. The method of claim 1, wherein a new cluster is created in response to one or more of said traffic measurements rising above a specified level.
14. The method of claim 1, wherein one of said clusters is eliminated in response to one or more of said traffic measurements falling below a specified level.
15. A method for cluster management in a network environment, comprising:
 - receiving a request from a node to change affiliation with said network environment;

determining if the affiliation change would result in an integer-squared number of nodes being affiliated with said environment; and

dispatching data to realize reclustering in said environment in the case where said determining yields an affirmative result.

16. The method of claim 15, wherein said network environment is a peer-to-peer environment.

17. The method of claim 15, wherein it is determined if the affiliation change would result in said integer-squared number of nodes being registered in said network environment.

18. The method of claim 15, wherein it is determined if the affiliation change would result in said integer-squared number of nodes being active in said network environment.

19. The method of claim 15, wherein the affiliation change is registration.

20. The method of claim 15, wherein the affiliation change is entry into active state.

21. The method of claim 15, wherein said reclustering comprises establishment of a new cluster in said network environment.

22. The method of claim 15, wherein said reclustering comprises elimination of an existing cluster in said network environment.

23. The method of claim 15, wherein said reclustering comprises transfer of one or more nodes from a first cluster in said network environment to a second cluster in said network environment.

24. A method for communications in a network environment, comprising:

receiving data at a node in said network environment, wherein said node is associated with a cluster in said network environment;

selecting from identification numbers associated with nodes in said network environment an identification number closest in value, in view of a specified polarity, to an identification number associated with said node; and

dispatching said data to a node associated with the selected identification number.

25. The method of claim 24, wherein said network environment is a peer-to-peer environment.

26. The method of claim 24, wherein the identification number associated with the node that received said data and the selected identification number are node identification numbers, and said node associated with the selected identification number is associated with said cluster.

27. The method of claim 24, wherein the identification number associated with the node that received said data and the selected identification number are cluster identification numbers, and node associated with the selected identification number is associated with a cluster other than the cluster with which the node that received said data is associated.

28. The method of claim 27, wherein said node associated with the selected identification

number is selected randomly from a plurality of nodes associated with the selected identification number.

29. The method of claim 26, wherein said data corresponds to an entity index update.

30. The method of claim 3, wherein said data corresponds to an entity query.

31. The method of claim 26, wherein said data corresponds to a node index update.

32. The method of claim 27, wherein said data corresponds to a node index update.

33. The method of claim 24, wherein said specified polarity indicates that the selected identification number be higher in value than the identification number associated with the node that received said data.

34. The method of claim 24, wherein said specified polarity indicates that the selected identification number be lower in value than the identification number associated with the node that received said data.

35. A system for cluster management in a network environment, comprising:

a memory having program code stored therein; and

a processor operatively connected to said memory for carrying out instructions in accordance with said stored program code;

wherein said program code, when executed by said processor, causes said processor to perform:

receiving one or more traffic measurements from one or more nodes associated with one or more clusters in said network environment;

determining, in accordance with the received traffic measurements, one or more reclustering operations to be performed in said network environment; and

dispatching data to realize said reclustering.

36. The system of claim 35, wherein said network environment is a peer-to-peer environment.

37. The system of claim 35, wherein said reclustering operations comprise creation of a new cluster.

38. The system of claim 35, wherein said reclustering operations comprise elimination of one of said clusters.

39. The system of claim 35, wherein said reclustering operations comprise transfer of one or more of said nodes among between one or more of said clusters.

40. The system of claim 35, wherein said traffic measurements are constantly taken.

41. The system of claim 35, wherein said traffic measurements are taken in response to a request for said measurements.

42. The system of claim 41, wherein said measurements are taken for a specified period of time.
43. The system of claim 35, wherein said traffic measurements comprise measurements corresponding to node index updates.
44. The system of claim 35, wherein said traffic measurements comprise measurements corresponding to entity index updates.
45. The system of claim 35, wherein said traffic measurements comprise measurements corresponding to entity queries.
46. The system of claim 35, wherein determining comprises determination of ideal cluster size.
47. The system of claim 35, wherein a new cluster is created in response to one or more of said traffic measurements rising above a specified level.
48. The system of claim 35, wherein one of said clusters is eliminated in response to one or more of said traffic measurements falling below a specified level.
49. A system for cluster management in a network environment, comprising:
- a memory having program code stored therein; and
 - a processor operatively connected to said memory for carrying out instructions in

accordance with said stored program code;

wherein said program code, when executed by said processor, causes said processor to perform:

receiving a request from a node to change affiliation with said network environment;

determining if the affiliation change would result in an integer-squared number of nodes being affiliated with said environment; and

dispatching data to realize reclustering in said environment in the case where said determining yields an affirmative result.

50. The system of claim 49, wherein said network environment is a peer-to-peer environment.

51. The system of claim 49, wherein it is determined if the affiliation change would result in said integer-squared number of nodes being registered in said network environment.

52. The system of claim 49, wherein it is determined if the affiliation change would result in said integer-squared number of nodes being active in said network environment.

53. The system of claim 49, wherein the affiliation change is registration.

54. The system of claim 49, wherein the affiliation change is entry into active state.

55. The system of claim 49, wherein said reclustering comprises establishment of a new cluster

in said network environment.

56. The system of claim 49, wherein said reclustering comprises elimination of an existing cluster in said network environment.

57. The system of claim 49, wherein said reclustering comprises transfer of one or more nodes from a first cluster in said network environment to a second cluster in said network environment.

58. A system for communications in a network environment, comprising:

a memory having program code stored therein; and

a processor operatively connected to said memory for carrying out instructions in accordance with said stored program code;

wherein said program code, when executed by said processor, causes said processor to perform:

receiving data at a node in said network environment, wherein said node is associated with a cluster in said network environment;

selecting from identification numbers associated with nodes in said network environment an identification number closest in value, in view of a specified polarity, to an identification number associated with said node; and

dispatching said data to a node associated with the selected identification number.

59. The system of claim 58, wherein said network environment is a peer-to-peer environment.

60. The system of claim 58, wherein the identification number associated with the node that received said data and the selected identification number are node identification numbers, and said node associated with the selected identification number is associated with said cluster.

61. The system of claim 58, wherein the identification number associated with the node that received said data and the selected identification number are cluster identification numbers, and node associated with the selected identification number is associated with a cluster other than the cluster with which the node that received said data is associated.

62. The system of claim 61, wherein said node associated with the selected identification number is selected randomly from a plurality of nodes associated with the selected identification number.

63. The system of claim 60, wherein said data corresponds to an entity index update.

64. The system of claim 37, wherein said data corresponds to an entity query.

65. The system of claim 60, wherein said data corresponds to a node index update.

66. The system of claim 61, wherein said data corresponds to a node index update.

67. The system of claim 58, wherein said specified polarity indicates that the selected identification number be higher in value than the identification number associated with the node that received said data.

68. The system of claim 58, wherein said specified polarity indicates that the selected identification number be lower in value than the identification number associated with the node that received said data.